

RATIONAL GENERATING SERIES FOR AFFINE PERMUTATION PATTERN AVOIDANCE

BRANT JONES

ABSTRACT. We consider the set of affine permutations that avoid a fixed permutation pattern. Crites has given a simple characterization for when this set is infinite. We find the generating series for this set using the Coxeter length statistic and prove that it can always be represented as a rational function. We also give a characterization of the patterns for which the coefficients of the generating series are periodic. The proofs exploit a new polyhedral encoding for the affine symmetric group.

1. INTRODUCTION

The affine symmetric group \tilde{S}_n is an infinite group that arises naturally in various geometric, combinatorial, and algebraic contexts. In this work, we are concerned with the enumeration of various subsets of this group. Since the group is infinite, we consider “refined” counts of elements based on the Coxeter length statistic $\ell(w)$ that describes the minimal number of generators needed to factor $w \in \tilde{S}_n$ in a certain standard group presentation of \tilde{S}_n . We will use the language of **generating series** to describe our results: For a given subset $S \subseteq \tilde{S}_n$, we form the series $\sum_{w \in S} x^{\ell(w)}$ using a formal variable x and attempt to find a closed form for this expression. The associated **enumerating sequence** is the sequence of coefficients which counts the number of elements of each given length. These are related; for example, the enumerating sequence is given by a linear constant-coefficient recurrence precisely when the generating series can be expressed as a rational function.

One of the first results in this direction is due to Bott [Bot56] who gave a general method to compute the Poincaré series that describes the Betti numbers for the associated compact Lie group. Combinatorially, this is the generating series by Coxeter length for the entire group $S = \tilde{S}_n$.

Theorem 1.1. [Bot56] *We have*

$$\sum_{w \in \tilde{S}_n} x^{\ell(w)} = \frac{(1+x)(1+x+x^2)(1+x+x^2+x^3) \cdots (1+x+x^2+x^3+\cdots+x^{n-1})}{(1-x)(1-x^2)(1-x^3) \cdots (1-x^{n-1})}$$

Although his motivation and proof were topological, it is relatively straightforward to give a combinatorial proof by induction using the (so-called “parabolic”) subgroups obtained from subsets of the standard generators (see [Hum90, (5.12)]). These subgroups turn out to be finite symmetric groups, for which the generating series is given by the numerator in Theorem 1.1. We give a new combinatorial proof for the denominator in Bott’s formula in Corollary 2.2.

Recently, Crites gave a natural extension of permutation pattern avoidance for the affine symmetric group as part of his thesis work with Sara Billey [BC12] to characterize the rationally smooth Schubert varieties of affine type A . In [Cri10], he also enumerated the number of affine permutations avoiding various fixed patterns, and proved the following remarkable structure theorem.

Theorem 1.2. [Cri10] *Let p be a finite permutation and $n \geq 2$. There exist only finitely many affine permutations of size n that avoid p if and only if p avoids the classical permutation pattern [321].*

Even when there are infinitely many affine permutations of size n that avoid a fixed pattern p , we can still consider the length generating series

$$F_{p,n}(x) := \sum_{\substack{w \in \tilde{S}_n \\ w \text{ avoids } p}} x^{\ell(w)}.$$

Such series first appeared in Hanusa and Jones' [HJ10] enumeration of the $[321]$ -avoiding affine permutations. It is shown there that the coefficients of the length generating series for $p = [321]$ are periodic. These $[321]$ -avoiding affine permutations are also known as the fully commutative elements of affine type A . More recently, Biagioli, Jouhet, and Nadeau [BJN13] have described the length generating series for fully commutative elements in other affine types, and they turn out to be periodic there as well. In fact, they propose the problem of determining which Coxeter groups have a periodic generating series associated to their subset of fully commutative elements. This would generalize Stembridge's classification [Ste96].

In this work, we consider the dual problem of classifying the periodic patterns within the affine symmetric group. While any generating series with periodic coefficients can be expressed as a reduced rational function with denominator $1 - x^d$, it is not obvious that the $F_{p,n}(x)$ series are even rational in general. One standard way to show that a counting problem is solved by a rational generating series is to produce a bijection to directed paths in a finite graph (or equivalently, words in a regular language). Stanley [Sta97] refers to this as the "transfer matrix method." In fact, Brink and Howlett have described a clever finite state automaton that recognizes a canonical reduced expression for each element of a fixed Coxeter group (see [BH93] or [BB05, Chapter 4]); Casselman has also contributed significantly to make their ideas practical for efficient implementation in software (see [Cas95], for example). We initially attempted to modify these constructions to filter the affine permutations based on pattern avoidance criteria. At this stage, however, it appears that pattern avoidance is not sufficiently related to the group structure for this approach to work in general.

Recently, we have turned instead to a set of ideas based on geometric convexity. Consider a rational polyhedron P defined as the set of solutions in \mathbb{R}^n to a set of linear inequalities with integral coefficients, and suppose that we would like to count the lattice points in $\mathbb{Z}^n \cap P$. To be more general, we consider the **encoding series**

$$\mathcal{F}_P(x_1, \dots, x_n) := \sum_{(z_1, \dots, z_n) \in \mathbb{Z}^n \cap P} x_1^{z_1} x_2^{z_2} \cdots x_n^{z_n}$$

for these points in the formal variables x_1, \dots, x_n . Brion's formula (see [BR07] or [Bar02]) states that this encoding series is simply the sum of the encoding series for each of the "tangent cones" formed by the rays emanating from a vertex of P . Moreover, it is straightforward to see (after using inclusion-exclusion if the cones are not simple) that the encoding series for these tangent cones are all rational, and so any generating series obtained by specializing the x_i will be rational also.

More precisely, we show in Section 2 how to coordinatize (the minimal length coset representatives of) \tilde{S}_n as the set of lattice points (z_1, \dots, z_{n-1}) in the nonnegative orthant $\mathbb{Z}_{\geq 0}^{n-1}$ with Coxeter length given by $\sum_{i=1}^{n-1} (n-i)z_i$. Enumerating these points recovers the denominator of Bott's formula.

However, it turns out that the subset of lattice points corresponding to the p -avoiding affine permutations, for a fixed pattern p , is not necessarily convex; see Figure 3(b). We then show that it is possible to decompose $\mathbb{Z}_{\geq 0}^{n-1}$ into a disjoint union of $(n-1)!$ shifted, dilated cones, each of the form

$$C_b^n := \{(t_1, 2t_2, \dots, (n-1)t_{n-1}) + (b_1, \dots, b_{n-1}) : t_i \in \mathbb{Z}_{\geq 0}\}.$$

If we restrict to each \mathcal{C}_b^n , then we can prove that the p -avoiding affine permutations do form a polyhedral set. In fact, we give explicit defining inequalities that include some additional coordinates for convenience, and then project to the t -coordinates that parameterize each \mathcal{C}_b^n . At the end of this process, we can apply Brion's formula to compute the enumerating series and conclude that it is rational.

Let us pause to mention that this construction seems likely to be useful in other contexts. For example, the Coxeter hyperplane arrangement of affine type A_{n-1} in \mathbb{R}^n is given by $x_i - x_j = k$ for $1 \leq i < j \leq n$ and $k \in \mathbb{Z}$. The complement of these hyperplanes in \mathbb{R}^n is a collection of regions. It turns out that these regions are in bijection with affine permutations, and so enumerating these regions using a statistic defined by counting the number of hyperplanes that separate a region from a fixed region at the origin results in the same generating series as Bott's formula. There is some recent interest [Arm13, FV10] in statistics and generating series for regions of the extended Shi arrangements (which are a subarrangements of this one), and affine pattern avoidance may be a useful tool for refining this geometric picture.

The generating series we have been considering also arise in certain lattice path enumeration problems; see [BDLPP01, BDLFP98, BJN13]. In fact, the enumeration for $p = [321]$ in [HJ10] used a recursive technique of Bousquet-Mélou [BM96] developed for this context involving q -Bessel functions that, while powerful, leaves the generating series in a form that is somewhat opaque. Our decomposition of the coordinate space for these objects into shifted dilated cones seems likely to offer some new insights into these types of recursive systems.

Once we know that our $F_{p,n}(x)$ generating series are rational, there are three possibilities for the sequence of coefficients: they must be eventually zero, eventually repeat, or are unbounded. The first case is characterized by Crites' theorem, and in Section 3, we begin to characterize the periodic patterns. We are aided by the fact that it suffices to characterize the periodic patterns in a single \mathcal{C}_b^n space, with $n = 3$. Stated in terms of classical permutation patterns, our result essentially requires p to avoid an infinite family of patterns from $S_7, S_8, S_{10}, S_{12}, S_{14}, \dots$; see Figure 5 and Theorem 3.11.

When p cannot be embedded into any element of \tilde{S}_n then the generating series $F_{p,n}(x)$ is simply given by Bott's formula, which is not periodic (unless $n = 2$). It remains an open problem to give a characterization in terms of p for when this occurs. We do not address this here although there are standard techniques from convex geometry that can be applied to the polyhedra we define for any particular pattern of interest.

There are many open directions in this area, for both undergraduate and professional researchers. Almost any of the classical problems associated with permutation patterns, such as classification of Wilf equivalence classes, pattern packing, or asymptotic behavior, could be posed in the affine setting; see [Bón12] for an introduction to these classical results. It would also be interesting to extend our geometric framework to study bivariate generating series of the form $\sum_{w \in S \subseteq \tilde{S}_n} x^{\ell(w)} y^n$. Moreover, modifying the geometric framework to handle multiple patterns would allow us to study the $\{[3412], [4231]\}$ class from [BC12] in detail.

2. POLYHEDRAL STRUCTURE

2.1. A polyhedral encoding of the affine symmetric group. An **affine permutation of size n** is a bijection $w : \mathbb{Z} \rightarrow \mathbb{Z}$ satisfying $w(i + n) = w(i) + n$ for all $i \in \mathbb{Z}$, and $w(1) + w(2) + \dots + w(n) = 1 + 2 + \dots + n$. We refer to the (infinite) image sequence $(\dots, w(-2), w(-1), w(0), w(1), \dots)$ of w as its **\mathbb{Z} -notation**. By the first property, we can completely specify an affine permutation by its **base-window** $[w(1), w(2), \dots, w(n)]$. When we do this, the \mathbb{Z} -notation is obtained by decomposing the image into **windows** of size n , where the i th window contains the entries of the base window with each value in the window shifted by in . (In this paper, we denote window boundaries with a $|$ symbol.)

The **affine symmetric group** \tilde{S}_n consists of all the affine permutations of size n , with composition of functions as the group operation. It follows directly from the definitions that $[w_1, w_2, \dots, w_n]$ is the base-window notation for an affine permutation if and only if $\sum_{i=1}^n w_i = \binom{n+1}{2}$ and the residues $(w_i \bmod n)$ are all distinct.

As a group, \tilde{S}_n is generated by the n adjacent transpositions of entries in the \mathbb{Z} -notation (where each transposition acts on all windows simultaneously). The minimal number of such transpositions into which w can be factored is known as the **Coxeter length** of w , denoted $\ell(w)$.

Given a permutation $p \in S_k$ and an affine permutation $w \in \tilde{S}_n$, we say that w **contains the pattern** p if there exist positions $i_1 < i_2 < \dots < i_k$ whose \mathbb{Z} -notation values $w(i_1), w(i_2), \dots, w(i_k)$ are in the same relative order as p_1, p_2, \dots, p_k . Note that these positions need not be restricted to the base window.

When the entries in the base-window notation for w are sorted increasingly, we call w a **minimal length coset representative**. (See [BB05, Hum90] for motivation and details.) We will denote the subset of minimal length coset representatives by $\tilde{S}_n^\circ \subseteq \tilde{S}_n$. Then each $w \in \tilde{S}_n^\circ$ corresponds to an **abacus diagram** as follows. Begin with an array having n columns and countably many rows. Label the entry in the i th row and j th column of the array by the integer $j + ni$, where $1 \leq j \leq n$. In figures, we will draw the rows increasingly up the page, and columns increasingly from left to right. Then these labels linearly order the entries of the array, which we refer to as **reading order**. We call the entries $\{1 + kn, 2 + kn, \dots, n + kn\}$ the k th **level** of the array. To create our diagram, we highlight certain entries in the array; such entries are called **beads** and will be circled in figures. Entries that are not beads will be called **gaps**. To encode w , we let the entries in the array corresponding to the base-window notation for w be beads, and we refer to these as the **defining beads**. To complete the diagram, we create beads at all of the entries below each defining bead, lying in the same column. All of the other entries in the diagram are gaps. We call this completed diagram the **abacus diagram for w** .




FIGURE 1. An abacus diagram for $w = [-12, -8, 2, 9, 13, 17]$ with $\bar{w} = (4, 10, 7, 4, 4)$ and $\hat{w} = (0, 3, 3, 2, 3)$.

Observe that the defining conditions on the base-window notation imply that the levels of the defining beads in an abacus diagram must sum to zero. We refer to this by saying that the abacus diagram must be **balanced**. Hence, the base window notation includes a redundant coordinate. To remedy this, we can represent any minimal length coset representative w by its **gap vector** $\hat{w} = (\hat{w}_1, \dots, \hat{w}_{n-1})$ where \hat{w}_i records the number of gaps between the i and $(i + 1)$ st defining beads in the abacus for w , ordered increasingly. Alternatively, we may specify w by its **delta vector**

$$\bar{w} = (w_2 - w_1, w_3 - w_2, \dots, w_n - w_{n-1}).$$

This vector records the number of entries (which may be beads or gaps) in the abacus diagram between each successive pair of defining beads.

Observe that *any* nonnegative integer vector is the gap vector for a unique abacus diagram. To see this, simply place the largest defining bead arbitrarily on the array, and then place each of the smaller defining beads with consecutive distances as prescribed by the given gap vector. To balance the abacus, subtract the sum of the levels of the defining beads from the position of each defining bead. The result will be the unique balanced abacus having the prescribed gap distances between consecutive defining beads.

Proposition 2.1. *The Coxeter length of w is given by $\ell(w) = \dot{w}_{n-1} + 2\dot{w}_{n-2} + 3\dot{w}_{n-3} + \cdots + (n-1)\dot{w}_1$.*

Proof. This is a “folklore” result that is sometimes stated in a slightly different form: To compute the Coxeter length of the element encoded by an abacus diagram, count the number of pairs (b, g) where b is a defining bead and g is a gap that precedes b in reading order.

Once we translate the action of \tilde{S}_n to the abacus, it is straightforward to prove this result by induction on $\ell(w)$; simply check that each length increasing adjacent transposition adds a single new (b, g) pair. \square

As a corollary to this development, we may view the (gap vectors of) minimal length coset representatives as lattice points in the nonnegative orthant, which is a prototype for our polyhedral encoding. When we enumerate these points with respect to the Coxeter length statistic, we recover the classical result of Bott for type \tilde{A} . This seems to be a new proof, and it is an open problem to give analogous proofs for the other affine Weyl groups.

Corollary 2.2. (Bott) *We have*

$$\sum_{w \in \tilde{S}_n^\circ} q^{\ell(w)} = \frac{1}{(1-q)(1-q^2) \cdots (1-q^{n-1})}.$$

Proof. By the development above, the encoding series for the gap vectors is

$$\sum_{g \in \mathbb{Z}_{\geq 0}^{n-1}} x_1^{g_1} x_2^{g_2} \cdots x_{n-1}^{g_{n-1}} = \frac{1}{(1-x_1)(1-x_2) \cdots (1-x_{n-1})}.$$

By Proposition 2.1, we can then obtain the length generating series by substituting q^{n-i} for x_i . This yields the result. \square


To uncover the polyhedral structure that will be useful in conjunction with patterns, we need a further refinement. We say that an abacus on n columns is **minimal** if its delta vector uses only entries between 1 and $n-1$. For example, the minimal abaci in $n=4$ are shown below in Figure 2.

Proposition 2.3. *There are $(n-1)!$ distinct minimal abaci on n columns.*

Proof. We argue by induction, the result being clear if $n=2$. Assume the formula holds for abaci on $n-1$ columns. To form a minimal abacus on n columns, we can start with a minimal abacus on $n-1$ columns, insert a new column containing a new largest defining bead in any of $n-1$ distinct positions, and rebalance the resulting n column abacus. Moreover, every minimal abacus on n columns arises this way. Hence, the formula holds by induction. \square

Given $w \in \tilde{S}_n^\circ$, we can project w to a minimal abacus by repeatedly removing multiples of n entries between consecutive defining beads and then rebalancing the diagram. We call the minimal abacus obtained in this way the **bias** of w . Equivalently, the bias \bar{b} of w is specified by its delta vector $\bar{b} = (\bar{w}_1 \bmod n, \bar{w}_2 \bmod n, \dots, \bar{w}_{n-1} \bmod n)$. Let BIAS_n denote the set of $(n-1)!$ possible biases on n columns.

Example 2.4. The bias of the abacus shown in Figure 2.1 is given by $\bar{b} = (4, 4, 1, 4, 4)$.

FIGURE 2. The $(n - 1)!$ different minimal abaci in $n = 4$

Lemma 2.5. *We can decompose the set of gap vectors into a disjoint union of shifted dilated cones*

$$\tilde{S}_n^\circ \cong \mathbb{Z}_{\geq 0}^{n-1} = \bigcup_{b \in \text{BIAS}_n} \mathcal{C}_b^n$$

where

$$\mathcal{C}_b^n := \{(t_1, 2t_2, \dots, (n-1)t_{n-1}) + (\mathring{b}_1, \dots, \mathring{b}_{n-1}) : t_i \in \mathbb{Z}_{\geq 0}\}.$$

Note that each element of \mathcal{C}_b^n is shifted by the same vector $(\mathring{b}_1, \dots, \mathring{b}_{n-1})$ that depends only on b . Hence we will refer to points in \mathcal{C}_b^n by their t -coordinates, (ab)using the notation $(t_1, \dots, t_{n-1}) \in \mathcal{C}_b^n$.

Proof. We claim that each gap vector $\mathring{w} = (\mathring{w}_1, \dots, \mathring{w}_{n-1}) \in \mathbb{Z}_{\geq 0}^{n-1}$ exists in precisely one of the \mathcal{C}_b^n sets. To see this, draw the abacus associated to the point w . Suppose the i and $(i+1)$ st defining beads have more than n entries between them. Then we delete one entire level of the array between them (and then renumber the remaining entries of the array). This maneuver removes i gaps from the coordinate \mathring{w}_i (since there will be $n-i$ beads on the level we remove), which is equivalent to removing 1 from coordinate t_i . Repeat this process until every consecutive pair of defining beads is separated by less than n entries, and then rebalance the abacus (by subtracting the sum of the levels of the defining beads from the position of each defining bead). By definition, the result will be one of the minimal abaci.

Moreover this process is reversible since we can recover w by starting with the minimal abacus, inserting levels as prescribed by the t_i coordinates, and rebalancing. Hence, the point of \mathcal{C}_b^n is unique. \square


Until now, we have focused on the minimal length coset representatives \tilde{S}_n° . From the length-additive parabolic decomposition in the theory of Coxeter groups, we have that the base-window notation of each $w \in \tilde{S}_n$ can be decomposed into a set of values together with a “sorting permutation” $v \in S_n$. The set of values is represented by some $u \in \tilde{S}_n^\circ$, and we have seen that these further decompose into subsets of elements having the same bias. The finite permutation v is the unique finite permutation having entries in the same relative order as the base-window notation for w . We call v the **flattening** of w , and it follows that $\ell(w) = \ell(u) + \ell(v)$.

Hence, we can extend our polyhedral embedding to \tilde{S}_n by simply taking $n!$ copies of the embedding for \tilde{S}_n° . Thus we let $\mathcal{C}_{b,v}^n$ denote the set of $w \in \tilde{S}_n$ whose bias is b and whose flattening is equal to v .

Corollary 2.6. *We have the disjoint union*

$$\tilde{S}_n \cong \bigcup_{\substack{b \in \text{BIAS}_n \\ v \in S_n}} c_{b,v}^n.$$

Example 2.7. In $n = 3$, we can draw the minimal length coset representatives as a set of lattice points in the plane. There are two minimal abaci, given by $\bar{a} = (1, 1)$ and $\bar{b} = (2, 2)$ with offsets given by $\dot{a} = (0, 0)$ and $\dot{b} = (0, 1)$, respectively. Then the gap vectors $\tilde{S}_3^\circ \cong \mathbb{Z}_{\geq 0}^2$ are a disjoint union of two cones, where the second coordinate has been dilated by 2 and the cones have been shifted by $(0, 0)$ and $(0, 1)$, respectively.



The entire affine symmetric group \tilde{S}_3 consists of six copies of this set of lattice points, one for each choice of flattening.

2.2. Patterns. Fix a permutation pattern $p \in S_k$, together with a bias b and flattening v . We will first explain how to characterize the elements of $C_{b,v}^n$ that contain an instance of p . Recall that an instance of the pattern p in the affine permutation w is a choice of k positions in the \mathbb{Z} -notation for w whose values have the same relative order as p . To coordinatize this, we consider two pieces of data associated to an instance: a strand assignment, and a window assignment.

Definition 2.8. Let the **strand assignment** of an instance be the function π assigning each entry of p to an entry of the base window of w , where $\pi(i) = j$ means that p_i is represented by some positional translation of the j th largest value of the base window (where $j = n$ represents the largest value).

The set of potential strand assignments for p is finite, consisting of all sequences of length k with values from $\{1, \dots, n\}$. In fact, it is not difficult to discover a further requirement for strand assignments.

Lemma 2.9. *Let π be a strand assignment for p . Then either every inversion in p must correspond to a strict inversion in π , or else the set of w containing p with strand assignment π is empty.*

Proof. Suppose $i < j$ and $p_i > p_j$. If $\pi(i) \leq \pi(j)$ then the elements representing p_i and p_j would necessarily be increasing in w , a contradiction. \square

Definition 2.10. The **window assignment** of an instance is the vector (c_1, \dots, c_{k-1}) where c_i is the number of positional window boundaries between the entries representing p_i and p_{i+1} in w . If p_i and p_{i+1} lie within the same window then we set $c_i = 0$.

Example 2.11. Consider the highlighted instance of $p = [24351]$ in the \mathbb{Z} -notation for $w = [-9, 4, 11]$:

$$(\dots - 15, -\mathbf{2}, \mathbf{5} | -12, \mathbf{1}, \mathbf{8} | -9, 4, 11 | -6, \mathbf{7}, 14 | -\mathbf{3}, 10, 17 | \dots)$$

Then, $\pi = [2, 3, 2, 2, 1]$ and $c = (0, 1, 2, 1)$.

It is clear that the strand assignment and the window assignment completely determine a pattern instance in w . Given an affine permutation w in \mathbb{Z} -notation, we can recover the t -coordinates from \mathcal{C}_b^n as follows.

Lemma 2.12. *Given $w \in \widetilde{S}_n^\circ$ let $j \geq 0$ be maximal such that $w(i+1) > w(i+jn)$. Then, t_i equals the number of window boundaries lying between $w(i+1)$ and $w(i+jn)$.*

For Example 2.11, we find $t_2 = 2$ since there are two window boundaries lying between entries 11 and 10. Similarly, $t_1 = 4$.

Proof. Work by induction starting from a minimal abacus. In a minimal abacus, there are no adjacent inversions between windows in \mathbb{Z} -notation and all the t_i are zero. Each time we add one to a t_i coordinate, we adjust w by adding n to each of the $n-i$ largest values in the base window, and then subtracting $n-i$ from each of the values in the base window (to rebalance). This places one new window boundary that is counted by the description in the statement, and preserves all of the other window boundaries. \square

We next characterize the t -coordinates of points in $\mathcal{C}_{b,v}^n$ that contain an instance of p with strand assignment π . To accomplish this, we highlight some data in (p, π) .

Definition 2.13. Given a pattern $p \in S_k$ and a strand assignment π for p , we say that an **upshift** is a pair $j < i$ such that $p_j = p_i + 1$ and $\pi(j) > \pi(i)$. A **downshift** is a pair $i < j$ such that $p_j = p_i + 1$ and $\pi(i) > \pi(j)$.

Example 2.14. Consider $(p, \pi) = ([24351], [2, 3, 2, 2, 1])$. The values 1 and 2 in p form an upshift that we denote (positionally) as $(1 < 5)$. Values 2 and 3 are both assigned to strand 2, so no shift takes place. Values 3 and 4 form the upshift $(2 < 3)$, and values 4 and 5 form the downshift $(2 < 4)$.

We are now in a position to prove our main result in this section.

Theorem 2.15. *For each (p, π) , the set*

$$\{t = (t_1, \dots, t_{n-1}) \in \mathcal{C}_{b,v}^n : t \text{ contains an instance of } p \text{ using strand assignment } \pi\}$$

consists of the integer points in a rational polyhedron.

Proof. We first prove that the set

$$\{(t, c) \in \mathcal{C}_{b,v}^n \times \mathbb{Z}_{\geq 0}^{k-1} : t \text{ contains an instance of } p \text{ using strand assignment } \pi \text{ and window assignment } c\}$$

consists of integer points in a rational polyhedron. Then, we can project onto the t -coordinates to obtain the result. More precisely, we will give a collection of integral linear inequalities that the (t, c) coordinates satisfy exactly when they describe a valid instance of p in the affine permutation corresponding to t . Then the fundamental Minkowski–Weyl Theorem for convex polyhedra allows us to describe this set as the Minkowski sum of a (bounded) polytope and a recession cone of rays. To perform the projection, we simply ignore the c coordinates in this latter description. (See [BR07] or [Zie95, Lecture 1] for an introduction to these ideas.)

Given a window assignment c , we imagine placing the values of p into an affine permutation w (whose values are to be determined) in increasing order. Whenever we place a larger value in a position to the right, on the same strand or higher, we impose no conditions on the t_i because the strands are necessarily increasing in w . However, if we place a larger value to the left, we must increase the strand and so this pair of consecutive values is an upshift. Also, if we place a larger value to the right on a lower strand, this pair of consecutive values is a downshift. These do impose conditions on the t_i .

By Lemma 2.12, we have that t_i represents the maximal number of positional window boundaries lying between (translates of) the i th and $(i+1)$ st largest elements of the base window that can form an inverted pair. Similarly, $t_i + t_{i+1} + \dots + t_{j-1} + \lfloor \bar{b}_{i,j} \rfloor$ represents the maximal number of positional window

boundaries lying between translates of the i th and j th largest elements of the base window that form an inverted pair. Here, $\lfloor \bar{b}_{i,j} \rfloor = \lfloor \frac{\bar{b}(i) + \bar{b}(i+1) + \dots + \bar{b}(j-1)}{n} \rfloor$ represents the (constant) contribution from the bias.

Therefore, when we have an upshift from strand $\pi(i)$ to strand $\pi(j)$, we must ensure that $t_{\pi(i)} + t_{\pi(i)+1} \dots + t_{\pi(j)-1} + \lfloor \bar{b}_{\pi(i), \pi(j)} \rfloor$ is large enough to ensure that the entries at positions j and i are inverted. These positions are separated by $c_j + c_{j+1} + \dots + c_{i-1}$ window boundaries.

Hence, for each upshift ($j < i$) we include an inequality of the form

$$t_{\pi(i)} + t_{\pi(i)+1} + \dots + t_{\pi(j)-1} + \lfloor \bar{b}_{\pi(i), \pi(j)} \rfloor \geq c_j + c_{j+1} + \dots + c_{i-1}.$$

Using similar reasoning for each downshift ($i < j$), we include an inequality of the form

$$t_{\pi(j)} + t_{\pi(j)+1} + \dots + t_{\pi(i)-1} + \lfloor \bar{b}_{\pi(j), \pi(i)} \rfloor \leq c_i + c_{i+1} + \dots + c_{j-1} - 1.$$

The choice of flattening enters in the initial conditions on c_i , which specify the minimal number of window boundaries between the i and $(i+1)$ st position in the pattern instance. Since these positions must be increasing, we include

$$c_i \geq \begin{cases} 0 & \text{if } v^{-1}(\pi(i)) < v^{-1}(\pi(i+1)) \\ 1 & \text{otherwise} \end{cases}$$

for each $1 \leq i < k$. The initial conditions $t_i \geq 0$ should also be included.

Our strategy to place p into the affine permutation will succeed if all of these linear inequalities are satisfied. If any are not satisfied, then we will have a pair of consecutive values from p whose representatives in the affine permutation do not faithfully represent the pattern. Hence, the integer points of this polyhedron form precisely the set given in the beginning of the argument. After projection, we obtain the result. \square

Example 2.16. The linear inequalities obtained for $(p, \pi) = ([24351], [2, 3, 2, 2, 1])$ with $\bar{b} = (0, 0)$ and $v = [123]$ are:

$$\begin{aligned} t_1 &\geq c_1 + c_2 + c_3 + c_4, & t_2 &\geq c_2, & t_2 &\leq c_2 + c_3 - 1, \\ c_1 &\geq 0, c_2 &\geq 1, c_3 &\geq 1, c_4 &\geq 1; & t_1 &\geq 0, t_2 &\geq 0. \end{aligned}$$

Definition 2.17. We will refer to the rational polyhedron constructed in the proof of Theorem 2.15 by $\mathcal{C}_{b,v}^n(p, \pi)$.

It turns out that the bias and flattening parameters do not change the polyhedra very much.

Lemma 2.18. Let b_0 be the bias given by $\bar{b} = (0, 0, \dots, 0)$ and v_0 be the identity permutation in S_n . Then for any other choice of $b \in \text{BIAS}_n$ and $v \in S_n$, we have that $\mathcal{C}_{b,v}^n(p, \pi)$ has the same set of infinite rays as $\mathcal{C}_{b_0,v_0}^n(p, \pi)$.

As a result, we often drop b and v from our notation, and let $\mathcal{C}^n(p, \pi) = \mathcal{C}_{b_0,v_0}^n(p, \pi)$.

Proof. Write the polyhedron $\mathcal{C}_{b,v}^n(p, \pi)$ as the solution set to a collection of linear inequalities. We can use a matrix A and multiplication by -1 to write this in a standard form $Ax \leq b$. It is well-known (and straightforward to verify) that changing b cannot change any of the infinite rays in the solution set. Since changing the bias or flattening parameters only alters the defining inequalities by a constant, and preserves all of the coefficients of the t_i and c_i , we obtain the result. \square

Corollary 2.19. For any permutation pattern p and any $n \geq 2$, the generating series

$$F_{p,n}(x) = \sum_{\substack{w \in \tilde{S}_n \\ w \text{ avoids } p}} x^{\ell(w)}$$

is rational. Equivalently, the coefficient sequence is generated by a linear constant-coefficient recurrence.

Proof. Using Brion's formula (see [BR07] or [Bar02]) together with inclusion-exclusion applied to Theorem 2.15, we can obtain a rational encoding series for the points of each $\bigcup_{\pi} \mathcal{C}_{b,v}^n(p, \pi)$. The subsequent union of these sets over all b and v are disjoint, so we can simply add the encoding series together. Then, we specialize the encoding series by setting t_i to $(x^i)^{n-i}$ for each $1 \leq i < n$. The first exponent dilates the lattice to recover the gap coordinates as in Lemma 2.5, and the second exponent comes from Proposition 2.1. Finally, we subtract the result from Bott's formula (which itself is rational) to enumerate the p -avoiding elements. \square

Let us turn to some examples in $n = 3$ where we can draw pictures.




FIGURE 3. (a) $(p, \pi) = ([321], [3, 2, 1])$; (b) $(p, \pi) = ([2431], [3, 3, 2, 1])$; (c) $(p, \pi) = ([24351], [2, 3, 2, 2, 1])$.

Example 2.20. In Figure 3 we have displayed some $\mathcal{C}_{b,v}^n(p, \pi)$. In each of the examples, we have $n = 3$, $v = [123]$, both biases are displayed superimposed, and (p, π) vary. We have also drawn some of the hyperplanes of constant Coxeter length from which the contributions to the rational generating series can be computed.

Observe that in Example (a) the counting sequence for the number of p -avoiding elements eventually stabilizes. In Example (b), we have a periodic sequence with period 2. Example (c) produces an unbounded counting sequence (although other strand assignments provide a ray in the y direction that is missing for this assignment; the full counting sequence for this p turns out to be periodic).

Warning 2.21. These polyhedra can be empty. For example, $p = [7, 1, 0, 4, 5, 2, 8, 10, 6, 9, 3]$ has only one strand assignment using 3 strands, and the corresponding $\mathcal{C}^n(p, \pi)$ polyhedron is empty.

There are some natural questions about these polyhedra to which we do not currently know the answer.

Question 2.22. *If we fix the bias and flattening parameters, is the union*

$$\bigcup_{\pi} \mathcal{C}_{b,v}^n(p, \pi)$$

over all strand assignments necessarily convex? (If so, this would dramatically simplify the computation of the encoding series.)

Question 2.23. *Given a pattern $p \in S_k$ with j strands, we certainly need $n \geq j$ in order to successfully embed p into \tilde{S}_n . By Warning 2.21, this inequality is sometimes strict. Is there a simple way to describe the minimal size of an affine permutation that contains a given pattern p ?*

3. PERIODIC PATTERNS

Let a_i denote the coefficients of the rational generating series $F_{p,n}(x)$ from Corollary 2.19. That is, a_i counts the number of affine permutations of fixed size n and length i that avoid the fixed pattern p . Since the a_i obey a recurrence, it follows that there are three possible types of behavior.

Definition 3.1. We say that a permutation pattern p is **finitely enumerated** if the a_i are eventually zero. We say p is **periodic** if the a_i eventually satisfy $a_i = a_{i-N}$ for some fixed N . Otherwise, we say that p is **unbounded**.

(To verify that this definition is etymologically sound, use the pigeonhole principle to show that whenever a_i is a bounded sequence that satisfies a recurrence using a fixed number of prior terms, then a_i is actually periodic.)

Crites' characterized the finitely enumerated patterns in [Cri10], and Hanusa–Jones gave the first example of a periodic pattern, $p = [321]$, in [HJ10]. Our goal in this section is to characterize all of the periodic patterns.

Note that the classification in Definition 3.1 depends only on the denominator of the generating series and so the contributions from each bias and flattening must each fall into the same case by Lemma 2.18. For this reason, it suffices to work with the enumerating sequence for $\bigcup_{\pi} \mathcal{C}^n(p, \pi)$ in this section.

Definition 3.2. Given $p \in S_k$, let m be the length of the longest decreasing subsequence of p . In this situation, we say that p **has m strands**.

We rephrase Crites' Theorem from [Cri10] as follows.

Theorem 3.3. (Crites) *In each n , the permutation pattern p is finitely enumerated if and only if p has fewer than 3 strands.*

The following result then shows that periodic patterns can only exist on three strands.

Proposition 3.4. *In each n , if p has four or more strands then p is unbounded.*

Proof. Consider Bott's formula for \tilde{S}_3° . The sequence of coefficients is unbounded, and the affine permutations in \tilde{S}_3° all avoid p (since the length of the longest decreasing subsequence in any of them is clearly 3 or less). When $n > 3$, we can embed $w \in \tilde{S}_3^\circ$ into \tilde{S}_n° by padding w with zeros on the left. This embedding is injective, the length of the longest decreasing subsequence in the image will be the same or smaller, and by Proposition 2.1 we do not change the Coxeter length. Hence, we obtain the result. \square

Lemma 3.5. *In each n , we have that p is periodic if and only if there exists a constant B such that $\bigcup_{\pi} \mathcal{C}^n(p, \pi)$ contains every point (t_1, \dots, t_{n-1}) that has two or more t_i coordinates larger than B .*

Proof. Since the enumerating sequences are generated by a recurrence, we have that p is periodic if and only if there exists an upper bound for the values of the sequence. Also, the condition in the statement for the t coordinates is true if and only if it is true for the corresponding w gap vector coordinates.

To prove the result, first suppose that at most one w_i coordinate (from the space $\mathbb{Z}_{\geq 0}^{n-1}$ of gap vectors) can become arbitrarily large when we avoid p . Then when we intersect with the hyperplane

$$w_{n-1} + 2w_{n-2} + 3w_{n-3} + \dots + (n-1)w_1 = i$$

of points with fixed Coxeter length i , for large i , the unbounded coordinate actually becomes determined; it must take up the “slack” in this equation for all of the bounded coordinates. As a result, we only have a bounded number of p -avoiding gap vectors, so the sequence is periodic.

Conversely, if there can be two unbounded gap vector coordinates when we avoid p , then one of these coordinates will be undetermined when we intersect with the hyperplane of fixed Coxeter length i , for sufficiently large i . Hence, the enumerating sequence is unbounded and so p is not periodic. \square

Corollary 3.6. *In $n = 3$, we have that p is periodic if and only if $\bigcup_{\pi} \mathcal{C}^3(p, \pi)$ contains infinite rays in the t_1 and t_2 directions.*

We now turn to classify the periodic patterns in $n = 3$. Eventually we show that these are the only periodic patterns for any n .

Definition 3.7. We say that (p, π) is **feasible** if $\bigcup_{\pi} \mathcal{C}^3(p, \pi)$ is nonempty.

While it remains an open problem to provide a (simple) combinatorial characterization for feasibility, there are standard techniques from convex geometry (such as Fourier–Motzkin elimination or Lenstra’s algorithm for integer programming [Sch86]) that may be used to address this question.

Definition 3.8. Let $p \in S_k$ with 3 strands, and let π be a strand assignment for p . Consider the diagram of (p, π) in which we represent p_i by a point (i, p_i) in the plane and label the point by its strand assignment $\pi(i)$.

We say that two elements of the second strand p_i and p_j are **linked below (above)** if there exists an element of the first (third, respectively) strand lying below and right (above and left, respectively) of both of them.

We say that two elements of the second strand are **chained below (above)** if there is a consecutive sequence of elements from the second strand between them that are linked below (above, respectively).

A **corner** of (p, π) consists of a triple (i, j, k) such that p_i and p_j are distinct elements of the second strand, and p_k is an element of the first or third strand that lies inside the square having p_i and p_j as diagonal vertices.

The corner is said to be **tight** if the elements p_i and p_j are chained below, or chained above.

Some tight corners that are chained below are shown schematically in Figures 4 and 5. Points drawn in the same row or column can be resolved to a permutation by either perturbation of the points. Thus, each picture encodes several classical permutation patterns.




FIGURE 4. Minimal tight corners

Lemma 3.9. *If (p, π) is feasible and contains a tight corner then t_1 or t_2 is not a ray of $\mathcal{C}^3(p, \pi)$.*

Proof. Suppose (p, π) has a tight corner. Without loss of generality, we may assume it is chained below as shown in the figures. Then, we claim that t_2 is not a ray. If it were, we could fix t_1 and increase t_2 arbitrarily. However, once t_1 is fixed, there is a maximum width for the strand 2 entries that are chained. Then we cannot increase t_2 past the distance limited by the strand 3 entry that is in the tight corner. \square

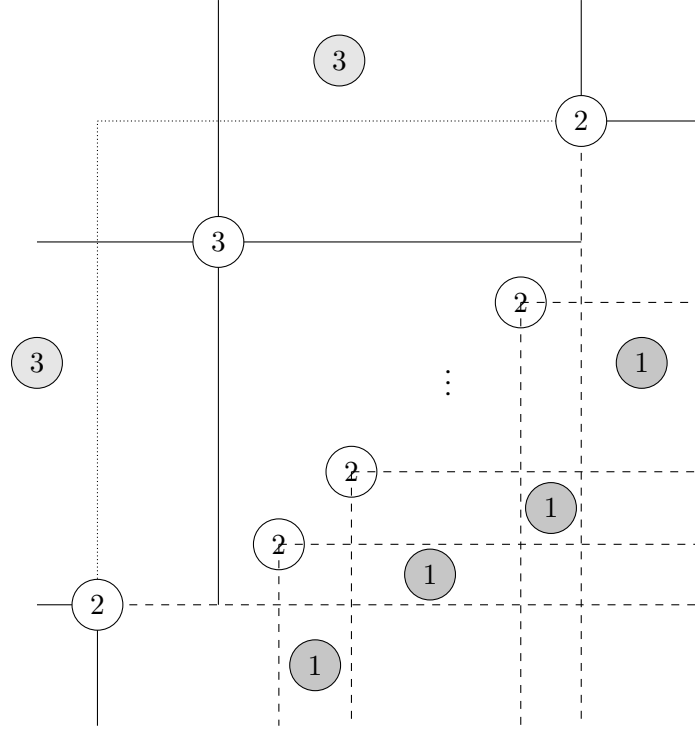


FIGURE 5. General form: tight corners

Lemma 3.10. *If (p, π) is feasible and does not contain a tight corner then both t_1 and t_2 are rays of $\mathcal{C}^3(p, \pi)$.*

Proof. We argue the contrapositive. Suppose, without loss of generality, that whenever t_1 is fixed there are only finitely many values for t_2 . If there were no corner of 2 entries enclosing a strand 3 entry, then we could separate strands 2 and 3, increasing t_2 arbitrarily. If the 2 entries defining the corner were not chained, then we could slide them along their strand and thereby increase t_2 arbitrarily. Therefore, we must have a tight corner. \square

Theorem 3.11. *The pattern p is periodic in \tilde{S}_3° if and only if there exists a strand assignment π that is feasible and does not contain a tight corner.*

Proof. First suppose there exists such a strand assignment. Then, Lemma 3.10 and Corollary 3.6 imply the result.

Next, suppose that no such strand assignment exists. If this is because no π is feasible, then p is not periodic since the enumeration is given by Bott's formula.

So suppose that every feasible π has a tight corner. We show that they all contain the same type of tight corner (i.e. are all chained above, or chained below). Fix some feasible π and consider the “supporting entries” shown in light gray in the figures. If these entries were not present in p , it would be possible to modify the strand assignment π to get rid of the tight corner, a contradiction.

Hence, the supporting entries must be present in p . But this implies that the strand assignments for the entries of the tight corner are forced in *every* strand assignment. Therefore, no π' can contain the ray that is missing due to the tight corner of π and Lemma 3.9. Thus, p is not periodic by Corollary 3.6. \square

Finally, we complete the periodic classification for $n > 3$.

Theorem 3.12. *Let $n > 3$. If p is not periodic in \tilde{S}_3° then p is not periodic in \tilde{S}_n° .*

Proof. Suppose for the sake of contradiction that p is not periodic in \tilde{S}_3° , but p is periodic in \tilde{S}_n° . Then, the values of the enumerating sequence for the p -avoiding elements of \tilde{S}_n° are bounded by some value B .

We embed points (t_1, t_2) from \tilde{S}_3° into \tilde{S}_n° by appending zeros on the left. After dilation, the length formula for these points is

$$\ell(0, 0, \dots, 0, t_1, t_2) = (n-1)t_2 + 2(n-2)t_1.$$

Hence, the points having fixed length k in \tilde{S}_n° satisfy

$$t_2 = -\frac{2(n-2)}{n-1}t_1 + \frac{1}{n-1}k.$$

Since p is not periodic in \tilde{S}_3° , we have that $\bigcup_\pi \mathcal{C}^3(p, \pi)$ does not contain rays in both the t_1 and t_2 directions. Since the points of $\bigcup_\pi \mathcal{C}^3(p, \pi)$ must lie in the nonnegative quadrant, this implies that some ray with positive slope separates the p -containing lattice points from the p -avoiding lattice points in $\mathbb{Z}_{\geq 0}^2$.

Because the embedded points having fixed length k in \tilde{S}_n° have negative rational slope, there must eventually be some large value of k for which we obtain more than B distinct points (t_1, t_2) such that:

- (1) $(0, \dots, 0, t_1, t_2)$ has length k in \tilde{S}_n° .
- (2) (t_1, t_2) avoids p in \tilde{S}_3° .

Since there are more than B of these embedded points, some of them must contain p . So suppose $t = (0, \dots, 0, t_1, t_2)$ contains p . Then there exists some strand assignment π for which t is feasible in $\mathcal{C}^n(p, \pi)$. Note that since the first $n-3$ coordinates are zero, we cannot have any upshifts involving strands $\{1, 2, \dots, n-2\}$. Therefore, we can form a strand assignment π' by changing all of these values to $n-2$, and the point t will still be feasible for $\mathcal{C}^n(p, \pi')$. But then the point t would also have been feasible for $\mathcal{C}^3(p, \pi'')$, where π'' is obtained from π' by sending $n-2$ to 1, $n-1$ to 2 and n to 3. This contradicts (2) above. \square

Theorem 3.13. *Let $n > 3$. If p is periodic in \tilde{S}_3° then p is periodic in \tilde{S}_n° .*

Proof. By Theorem 3.11, there exists a feasible strand assignment π whose solutions include both t_1 and t_2 as rays, so the solutions include all points having both coordinates larger than some B' . Let $t' = (t_1, \dots, t_{n-1}) \in \bigcup_\pi \mathcal{C}^n(p, \pi)$ with two coordinates larger than B' , say t_i and t_j .

Given π , we can form π' from π by preserving the images equal to 1, replacing the images equal to 2 by j , and replacing the images equal to 3 by n . This has the effect of replacing every instance of t_1 by $t_1 + \dots + t_{j-1}$ and every instance of t_2 by $t_j + \dots + t_{n-1}$ in the upshift and downshift inequalities from the proof of Theorem 2.15.

Hence, because the solutions for π include all points with $t_1 > B'$ and $t_2 > B'$, the solutions for π' include all points

$$t_1 + \dots + t_i + \dots + t_{j-1} > B'$$

and

$$t_j + \dots + t_{n-1} > B'.$$

Therefore, t' is a feasible point for π' .

Thus, \tilde{S}_n° satisfies Lemma 3.5, so is periodic. \square

ACKNOWLEDGEMENTS

We are grateful to Chris Hanusa, Edwin O'Shea, Len Van Wyk, and Alex Woo for helpful conversations about this work. Some problems related to the topic of this paper were investigated in undergraduate research programs mentored by the author at James Madison University, and we thank students Traymon Beavers, Karina Bekova, Abigail Liskey, and Ryan Stees for their contributions.

REFERENCES

- [Arm13] Drew Armstrong. Hyperplane arrangements and diagonal harmonics. *J. Comb.*, 4(2):157–190, 2013.
- [Bar02] Alexander Barvinok. *A course in convexity*, volume 54 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2002.
- [BB05] Anders Björner and Francesco Brenti. *Combinatorics of Coxeter groups*, volume 231 of *Graduate Texts in Mathematics*. Springer, New York, 2005.
- [BC12] Sara Billey and Andrew Crites. Pattern characterization of rationally smooth affine Schubert varieties of type A . *J. Algebra*, 361:107–133, 2012.
- [BDLFP98] E. Barucci, A. Del Lungo, J. M. Fédou, and R. Pinzani. Steep polyominoes, q -Motzkin numbers and q -Bessel functions. *Discrete Math.*, 189(1-3):21–42, 1998.
- [BDLPP01] E. Barucci, A. Del Lungo, E. Pergola, and R. Pinzani. Some permutations with forbidden subsequences and their inversion number. *Discrete Math.*, 234(1-3):1–15, 2001.
- [BH93] Brigitte Brink and Robert B. Howlett. A finiteness property and an automatic structure for Coxeter groups. *Math. Ann.*, 296(1):179–190, 1993.
- [BJN13] Riccardo Biagioli, Frédéric Jouhet, and Philippe Nadeau. Fully commutative elements and lattice walks. In *25th International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2013)*, Discrete Mathematics & Theoretical Computer Science Proceedings, AS. The Association. Discrete Mathematics & Theoretical Computer Science (DMTCS), Nancy, 2013.
- [BM96] Mireille Bousquet-Mélou. A method for the enumeration of various classes of column-convex polygons. *Discrete Math.*, 154(1-3):1–25, 1996.
- [Bón12] Miklós Bóna. *Combinatorics of permutations*. Discrete Mathematics and its Applications (Boca Raton). CRC Press, Boca Raton, FL, second edition, 2012. With a foreword by Richard Stanley.
- [Bot56] Raoul Bott. An application of the Morse theory to the topology of Lie-groups. *Bull. Soc. Math. France*, 84:251–281, 1956.
- [BR07] Matthias Beck and Sinai Robins. *Computing the continuous discretely*. Undergraduate Texts in Mathematics. Springer, New York, 2007. Integer-point enumeration in polyhedra.
- [Cas95] W. A. Casselman. Automata to perform basic calculations in Coxeter groups. In *Representations of groups (Banff, AB, 1994)*, volume 16 of *CMS Conf. Proc.*, pages 35–58. Amer. Math. Soc., Providence, RI, 1995.
- [Cri10] Andrew Crites. Enumerating pattern avoidance for affine permutations. *Electron. J. Combin.*, 17(1):Research Paper 127, 13, 2010.
- [FV10] Susanna Fishel and Monica Vazirani. A bijection between dominant Shi regions and core partitions. *European J. Combin.*, 31(8):2087–2101, 2010.
- [HJ10] Christopher R. H. Hanusa and Brant C. Jones. The enumeration of fully commutative affine permutations. *European J. Combin.*, 31(5):1342–1359, 2010.
- [Hum90] James E. Humphreys. *Reflection groups and Coxeter groups*, volume 29 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1990.
- [Sch86] Alexander Schrijver. *Theory of linear and integer programming*. Wiley-Interscience Series in Discrete Mathematics. John Wiley & Sons, Ltd., Chichester, 1986. A Wiley-Interscience Publication.
- [Sta97] Richard P. Stanley. *Enumerative combinatorics. Vol. 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1997. With a foreword by Gian-Carlo Rota, Corrected reprint of the 1986 original.
- [Ste96] John R. Stembridge. On the fully commutative elements of Coxeter groups. *J. Algebraic Combin.*, 5(4):353–385, 1996.
- [Zie95] Günter M. Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.

DEPARTMENT OF MATHEMATICS AND STATISTICS, MSC 1911, JAMES MADISON UNIVERSITY, HARRISONBURG, VA 22807

E-mail address: jones3bc@jmu.edu

URL: <http://educ.jmu.edu/~jones3bc/>